arXiv:1603.05143v2 [cond-mat.mtrl-sci] 18 Apr 2016

Robust Structural Identification via Polyhedral Template Matching

Peter Mahler Larsen, Søren Schmidt, Jakob Schiøtz

Department of Physics, Technical University of Denmark, 2800 Kgs. Lyngby, Denmark

Abstract. Successful scientific applications of large-scale molecular dynamics often rely on automated methods for identifying the local crystalline structure of condensed phases. Many existing methods for structural identification, such as Common Neighbour Analysis, rely on interatomic distances (or thresholds thereof) to classify atomic structure. As a consequence they are sensitive to strain and thermal displacements, and preprocessing such as quenching or temporal averaging of the atomic positions is necessary to provide reliable identifications. We propose a new method, Polyhedral Template Matching (PTM), which classifies structures according to the topology of the local atomic environment, without any ambiguity in the classification, and with greater reliability than e.g. Common Neighbour Analysis in the presence of thermal fluctuations. We demonstrate that the method can reliably be used to identify structures even in simulations near the melting point, and that it can identify the most common ordered alloy structures as well. In addition, the method makes it easy to identify the local lattice orientation in polycrystalline samples, and to calculate the local strain tensor. An implementation is made available under a Free and Open Source Software license.

1. Introduction

Often, the most challenging part of a Molecular Dynamics (MD) simulation is analysing the large amounts of data generated. For simulations within condensed matter physics, algorithms for automatic determination of local structure are often necessary, for example to identify crystalline phases, defects in the crystal structure, or structural motifs in non-crystalline samples. One method is to use the energy of the individual atoms and identify defect atoms by a threshold value of the energy. Another possibility is the centro-symmetry parameter, which is zero in many crystalline phases, but takes a non-zero value at defects such as dislocation cores [1]. Both methods are reasonably effective at identifying atoms near defects if the temperature is not too high, but do not provide additional information about the local structure.

A very popular analysis method is Common Neighbour Analysis (CNA) [2, 3], which classifies bonds between atoms according the local bonding structure, and uses this to assign a local crystalline structure to the atoms. The CNA has successfully been used to identify dislocations and grain boundaries in deforming polycrystalline metals [4, 5, 6], local ordering in amorphous phases [7, 8], and the competition between crystalline and icosahedral order in nanoparticles [9, 10].

The CNA method analyses the bonds between common neighbours of two bonded atoms. It relies heavily on the concept of two atoms being "bonded" or "neighbours", and thus needs a strict definition of this; typically in the form of a cut-off distance defining if two atoms are neighbours. Such a cut-off distance is, by necessity, somewhat arbitrary. Furthermore, thermal vibrations, coexistence of various phases, and fluctuations in the local density may all influence the result and either introduce noise in the analysis, or make it impossible to find a cutoff parameter useful for the entire system. For these reasons, Stukowski introduced the Adaptive Common Neighbour Analysis (ACNA) [11], where the cutoff distance is picked automatically, and individually, for each atom. While this makes the method significantly more robust, the ACNA still suffers from noise introduced by thermal vibrations, causing bonds to be sporadically broken or formed.

Recently, Lazar *et al.* have introduced a method classifying the local structure by the topology of the Voronoi cell surrounding the atom [12]. While this makes the method more robust to thermal vibrations than the (A)CNA, the method has difficulty distinguishing between the face-centered cubic (fcc) and hexagonal closed packed (hcp) structures.

In this paper, we introduce the Polyhedral Template Matching (PTM), an approach which is similar in spirit to the method of Lazar *et al.*. The gist of the method is that the convex hull formed by the set of neighbouring atoms describes the local structure around an atom. The convex hull is represented as a planar graph, and this graph is then used to classify the structure. As this method looks at a fixed number of neighbouring atoms around the atom being analysed, and as it does not employ a concept of bonds between these atoms, it is less sensitive to thermal fluctuations. In addition, the method assigns

a well-defined order to the neighbours of an atom, making it much easier to define a local orientation or a local elastic strain without referring to an initial unstrained structure. An implementation of the method is available online [13].

The structure of the paper is as follows: in section 2, we review how structures can easily be evaluated against a template structure, given that mapping between the neighbouring atoms and the template atoms exists. In section 3 we describe two methods for selecting the correct neighbouring atoms. In section 4 we show how the convex hull of the neighbouring atoms can be used to define a low number of candidate templates and mappings, and present the resulting algorithm for structure identification. In section 5, we benchmark the algorithm. In section 6 we extend it to ordered alloys. Finally, in sections 7 and 8, we illustrate how the method can provide information about local lattice orientation and local elastic strain, the former at no additional computational cost.

2. Template Matching

In this section we will describe how template matching can be used to choose the best structural match, given that a point-to-point correspondence between simulated and template structures exists. The process of finding the point-to-point correspondences is described in the next section. A commonly used measure of similarity between two point sets is the Root-Mean-Square Deviation (RMSD). Given two sets of points $\mathbf{v}$ and $\mathbf{w}$, the RMSD is defined as:

$$\mathrm{RMSD}(\mathbf{v},\mathbf{w}) = \sqrt{\frac{1}{N}\sum_{i=1}^{N} ||\vec{v_i} - \vec{w_i}||_2^2} \tag{1}$$

The superposition problem is that of finding a translation and a rotation of $\mathbf{w}$ and a scaling of $\mathbf{v}$ which minimizes the RMSD. This is equivalent to Horn's scale-asymmetric formulation [14]. It can be shown that the optimal translation is equivalent to bringing the barycentre of each point set to the origin. The optimal rotation and scaling are given by:

$$\mathrm{RMSD}(\mathbf{v},\mathbf{w}) = \min_{s,\mathbf{Q}} \sqrt{\frac{1}{N}\sum_{i=1}^{N} \left|\left| s[\vec{v_i} - \overline{\mathbf{v}}] - (\mathbf{Q}[\vec{w_i} - \overline{\mathbf{w}}]^T)^T \right|\right|_2^2} \tag{2}$$

where $\mathbf{Q}$ is a right-handed orthogonal matrix, $\overline{\mathbf{v}} = \frac{1}{N}\sum_{i=1}^{N}\vec{v_i}$ and $\overline{\mathbf{w}} = \frac{1}{N}\sum_{i=1}^{N}\vec{w_i}$ are the barycentres of $\mathbf{v}$ and $\mathbf{w}$ respectively, and s is the optimal scaling of $\mathbf{v}$. Finding $\mathbf{Q}$ is a well studied problem with many different solution methods; Theobald [15] provides a good exposition of the problem. Horn [14] describes a solution for finding s and shows that $\mathbf{Q}$ is independent of s. We can make the RMSD scale invariant by scaling $\mathbf{w}$ such

Structure	Neighbours Required
SC	6
FCC	12
HCP	12
ICO	12
BCC	14

Table 1. Number of neighbouring atoms required to identify structures.

that the mean distance of each point from the origin (after translation) is 1:

$$\mathrm{RMSD}(\mathbf{v},\mathbf{w}) = \min_{s,\mathbf{Q}} \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left|\left|s[\vec{v_i}-\overline{\mathbf{v}}]-\frac{1}{l(\mathbf{w})}\left(\mathbf{Q}[\vec{w_i}-\overline{\mathbf{w}}]^T\right)^T\right|\right|_2^2} \tag{3}$$

where:

$$l(\mathbf{w}) = \frac{1}{N}\sum_{i=1}^{N}||\vec{w_i}-\overline{\mathbf{w}}|| \tag{4}$$

Suppose now that we want to determine the structure of a central atom and its neighbours. We wish to determine if it has simple cubic (SC), face-centred cubic (FCC), hexagonal close packed (HCP), icosahedral (ICO) or body-centred cubic (BCC) structure. For SC, FCC, HCP and ICO the positions of the first shell of neighbouring atoms are sufficient to identify the structure. For BCC the first two shells are required. Table 1 shows the number of atoms required for each structure. Correct identifying the shell to which a neighbouring atom belongs is nontrivial at high temperatures; this is discussed in section 3. Given a set of reference templates corresponding to the atom positions of the aforementioned structures, the template which best matches an atom and its neighbours is the template which minimizes the RMSD after superposition. In the above formulation, $\mathbf{v}$ contains the positions of the central atom and its neighbours and $\mathbf{w}$ contains those of the template. A scale-invariant RMSD serves two purposes; it avoids preferential weighting of smaller templates and avoids the need for selecting bond lengths.

As we have demonstrated here, the task of structural identification would be simple if the optimal point-to-point correspondences were known. Clearly, a brute force approach of testing all possible permutations of the neighbours is computationally infeasible. Fast determination of point-to-point correspondences is the main contribution of this work; the algorithm for doing this is developed in section 4.

3. Neighbouring Atom Shell Identification

In order to use the template matching approach, we must correctly identify the nearest neighbours of each atom. In a perfect lattice, all atoms in the same neighbour shell lie at the same distance from a central atom, by definition. At low temperatures, thermal

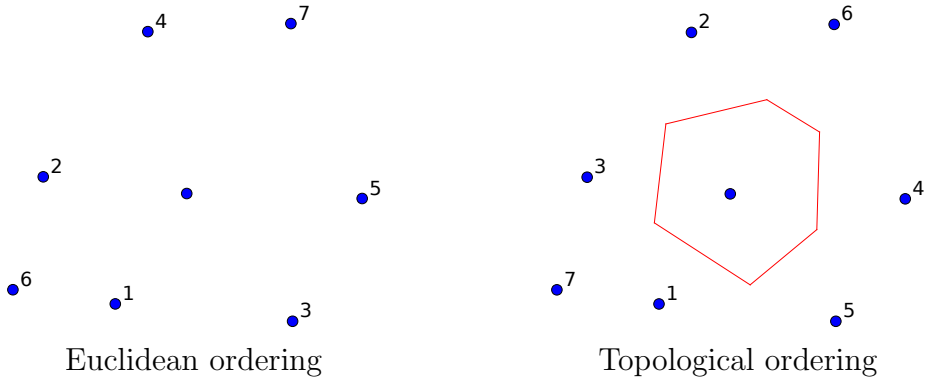

Figure 1. Neighbour atoms in a 2D hexagonal lattice ordered by Euclidean distance from the central atom (left) and by lexicographic ordering of the area of the polygonal interface between Voronoi cells and Euclidean distance (right). A 2D example is used here for ease of illustration. In $\mathbb{R}^2$, the 'polygonal' interfaces between Voronoi cells are line segments, and the interfacial area of each is simply the length.

displacements are small enough such that the Euclidean distances from the central atom are sufficient to identify the shell numbers of neighbouring atoms. At high temperatures, the thermal displacements are large enough that an atom in the second neighbour shell can be closer to the central atom than one in the first neighbour shell. In this case, rather than using distance to order the neighbouring atoms, we can order them using the areas of the bounding polygons of the Voronoi cell of the central atom. Given a discrete collection of points $\mathbf{P} = \{\vec{p_1}, \vec{p_2}, \ldots, \vec{p_N}\} \in \mathbb{R}^3$, the Voronoi cell of a point p_i consists of all points in $\mathbb{R}^3$ which are at least as close to p_i as to any other p_j. The boundary of a Voronoi cell can be defined by a set of polygons, each of which defines the interface to an adjacent Voronoi cell. Since we also wish to order neighbouring atoms whose Voronoi cells are not adjacent to that of the central atom, we use a lexicographical ordering of the interfacial areas *and* the distances:

$$\left(A_i, \frac{1}{d_i}\right) \geq \left(A_j, \frac{1}{d_j}\right) \qquad \forall i < j \tag{5}$$

where A_i is the area of the polygonal interface between the Voronoi cells of the central atom and a neighbouring atom, and d_i is the Euclidean distance between them. In the case where the Voronoi cell of a neighbouring atom is not adjacent to that of the central atom, we assign an area of $A_i = 0$. We denote this ordering as the 'topological' ordering, which is demonstrated for a 2D example in Figure 1.

4. Convex Hulls

We have described how template matching can be used to find the best structural match, given a known point-to-point correspondence. In this section we will describe how convex hull graphs can be used find all possible point-to-point correspondences in an efficient manner. A convex hull of a set of vertices is the smallest convex set which

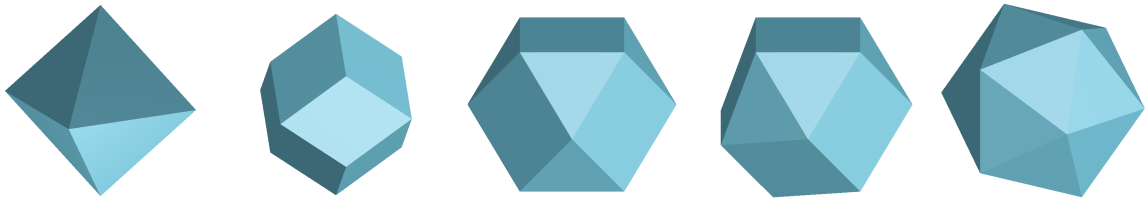

Figure 2. Convex hulls of the nearest neighbours of five different structures. For the BCC structure, the first two shells of nearest neighbours are included.

contains all the vertices. In $\mathbb{R}^3$ it can be described by a set of bounding planes. Figure 2 shows the convex hulls of the five different structural types we wish to match. The convex hull is the polyhedron formed by the nearest neighbours of a central atom.

Identification of the SC, FCC, HCP and ICO structures requires only the vertices of the first neighbour shell, which form a convex set. Two neighbour shells are required for identification of the BCC structure but, remarkably, these vertices also form a convex set. A stronger requirement of a convex hull is that it must consist only of simplicial facets, which in $\mathbb{R}^3$ are triangles; this is a *triangulation* of the surface of the convex hull. We can furthermore require that the triangulations fulfil the empty-sphere condition, that is, no vertex is contained inside the circumcircle of another simplicial facet. This is known as a *Delaunay* triangulation. Multiple Delaunay triangulations can exist for convex hulls with more than three coplanar points. Let us first consider the SC and ICO cases. The convex hulls of these structures consist only of triangular facets and the convex hull is therefore a unique triangulation. The FCC and HCP convex hulls consist of both equilateral triangular facets and perfect square facets. A perfect square has two equally valid triangulations which means the convex hull triangulation is not unique. Furthermore, the small atomic displacements can change the triangulation of the square facets. Figure 3 shows four different, equally valid, triangulations of the FCC convex hull.

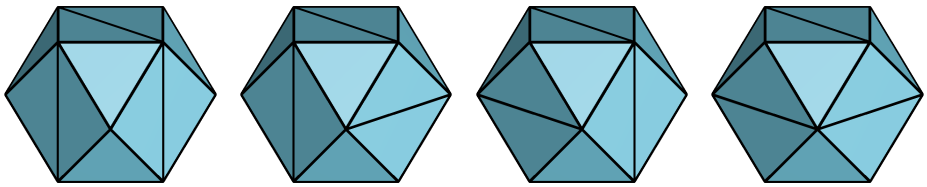

Figure 3. Four different triangulations of the FCC convex hull. The convex hull contains six square facets, which gives a total of 2^6 different triangulations.

The BCC convex hull consists only of rhombus facets. Whilst an unperturbed rhombus facet has a unique Delaunay triangulation, small changes in vertex positions can change the triangulation. We will consider both triangulations of each rhombus facet, since this will make our algorithm robust against relatively large atomic displacements.

We can represent the convex hull triangulations in graph form; vertices and edges in the convex hull triangulation correspond to vertices and edges in the graph. By representing a convex hull as a graph, we move from a metric space to a purely topological space. Steinitz's theorem [16] states that the skeleton of a three-dimensional convex polyhedron is a tri-connected planar graph. Conversely, any tri-connected planar graph has two embeddings in $\mathbb{R}^3$, i.e. it corresponds to two polyhedra that are mirror images of each other [17]. Figure 4 shows the planar graph representations of the convex hulls shown in Figure 2. Since the FCC, HCP and BCC convex hull triangulations are not unique, neither are their planar graph representations.

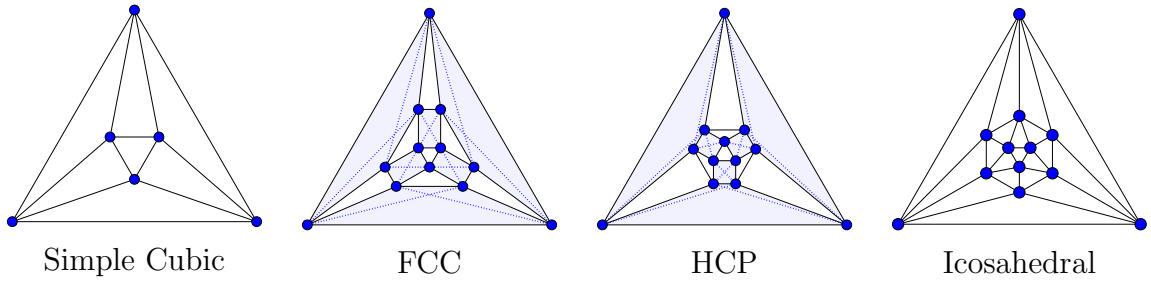

Figure 4. Convex hulls in planar graph form. The simple cubic and icosahedral graphs are unique, and are shown here as Tutte embeddings [18]. The FCC and HCP convex hulls do not have unique triangulations. The regions of the planar graphs which correspond to the square facets have been shaded. The dotted lines represent the two possible triangulations of each square facet. The BCC convex hull has no triangular facets, and cannot therefore be represented in the above manner.

We will use graph isomorphism to determine possible structure matches. Two graphs $G = (V, E)$ and $H = (W, F)$ (where V and W are the graph vertices and E and F are the graph edges) are said to be *isomorphic* if there exists a bijective function $f : V \to W$ such that for all $v, w \in V : \ \{v, w\} \in E \Leftrightarrow \{f(v), f(w)\} \in F$. An *automorphism* is an isomorphism of G to itself. Many practical algorithms determine graph isomorphism by computing the *canonical form* of each graph. The canonical form of a graph G is a uniquely defined automorphism of G such that any two graphs are isomorphic if and only if their canonical forms are identical. Weinberg [19] provides a simple method for defining the canonical forms and finding the automorphisms of tri-connected planar graphs (a more readable description is given by Kukluk et al. [20]). Weinberg's algorithm finds the canonical form by investigating both possible embeddings of a planar graph. Since we are dealing with the graphs of polyhedra in $\mathbb{R}^3$, an embedding already exists. We have therefore adapted the algorithm to investigate only a single embedding. By doing so, we restrict graph isomorphism to *orientation-preserving isomorphism*. Orientation-preserving isomorphism extends the above definition of graph isomorphism (we use the notation of Brinkmann et al. [21]): two *embedded* graphs are said to be orientation-preservingly isomorphic if there exists a bijective function $f_e : E \to F$ such that if $\{e_1, e_2, \ldots..., e_k\}$ is the set of edges incident with a vertex $v \in V$, in clockwise order, then $\{f_e(e_1), f_e(e_2), \ldots, f_e(e_k)\}$ is the set of edges incident with the vertex $f(v) \in W$, in clockwise order. An *orientation-preserving automorphism*

is an orientation-preserving isomorphism of G to itself. Henceforth, when we refer to isomorphisms and automorphisms, the orientation preserving variants are implied.

Using the planar graphs we can rule out impossible structure matches: in order for a set of points to match a reference template, the planar graphs of their convex hulls must be isomorphic. Unfortunately, the converse is not true; due to the multiple triangulations of the FCC and HCP convex hulls, the planar graphs are insufficient to identify structures uniquely. There are triangulations of FCC and HCP convex hulls whose planar graphs are isomorphic, and triangulations of both which are isomorphic to the icosahedral planar graph. As such, using planar graph isomorphism alone would sometimes lead to more than one matching structure. Nevertheless, the planar graph representation gives us a set of point-to-point correspondences to investigate, one of which is the optimal correspondence. We define the optimal correspondence as the one which minimizes the RMSD between a set of points and a reference template after superposition (as defined in Equation (3)). The total number of possible correspondences to investigate for each reference template is the product of the number triangulations of the convex hull, the number of automorphisms of each planar graph, and the two embeddings of each planar graph in $\mathbb{R}^3$. In practice, both the number of triangulations and automorphisms required is greatly reduced by the symmetries of each reference template (24-fold symmetry for simple cubic, BCC and FCC, 6-fold symmetry for HCP, 60-fold symmetry for icosahedral); the algorithm for generating the set of symmetrically inequivalent triangulations and automorphisms is described in Appendix A.

The structural identification process is described in Algorithm 1. An outline of the algorithm is as follows: for each atom, we loop over the possible structures to identify. The convex hull formed by the neighbouring atoms is calculated (the number of neighbouring atoms depends on the candidate structure), and the canonical form of the corresponding graph is found. We then loop over all possible triangulations of the reference structure (skipping triangulations that are symmetrically equivalent), and calculate the corresponding canonical form. If the canonical form of the reference structure is identical to the canonical form of the actual complex hull, we have a *possible* structural match. The quality of that match is then tested. This test is done by iterating over all symmetrically inequivalent automorphisms of the candidate graph, generating all possible structural templates which are then tested against the actual positions of the neighbouring atoms by optimizing the RMSD in Equation (3). The structure with the lowest RMSD is identified, and the algorithm returns the RMSD value, the corresponding orientation and the kind of structure identified.

5. Results

5.1. Copper Precipitate Benchmark

To measure the capabilities of the method, we will first make a comparison with two existing structural identification methods, Adaptive Common Neighbour Analysis

Algorithm 1 Pseudocode for determining local structure around a single atom. As argument, the procedure takes the positions of the atom and its neighbours, sorted in distance from the central atom. The algorithm returns the RMSD of the match found (or infinity if no match), the rotation matrix corresponding to the match, and the structure.

procedure DETERMINESTRUCTURE($\mathbf{A}$)
▷ $\mathbf{A}$ is the set of positions of central atom and its nearest neighbours
▷ $\mathbf{A}$ is ordered using Euclidean or topological ordering (c.f. sec. 3)
$r^* := \infty$ ▷ RMSD of best match
$\mathbf{Q}^* := \mathbf{1}$ ▷ Unit rotation
$\mathbf{S}^* :=$ disordered ▷ No structure identified (yet)
for $\mathbf{S} \in \{$SC, FCC, HCP, ICO, BCC$\}$ **do**
$\mathbf{W} :=$ ReferenceTemplate($\mathbf{S}$)
$\mathbf{U} = \{\vec{a_j} \in \mathbf{A} \mid j \leq |\mathbf{W}|\}$ ▷ Select innermost atoms
$\mathbf{C} :=$ Conv($\mathbf{U}$) ▷ Calculate convex hull of $\mathbf{U}$
if $\vec{a}_1 \notin \mathbf{C}$ **then** ▷ Convex hull must not contain central atom
$\mathbf{G} :=$ CanonicalForm(Graph($\mathbf{C}$))
for each triangulation $\mathbf{T}_i$ of Conv($\mathbf{W}$) **do**
$\mathbf{G}_{\text{ref}} :=$ CanonicalForm(Graph($\mathbf{T}_i$))
if $\mathbf{G} = \mathbf{G}_{\text{ref}}$ **then** ▷ Test graph isomorphism
for each automorphism $\mathbf{A}_j$ of $\mathbf{G}$ **do**
$\mathbf{U}' := \mathbf{U} \circ \mathbf{A}_j$ ▷ Permute by automorphism
$\{r, \mathbf{Q}\} :=$ RMSD($\mathbf{U}', \mathbf{W}$) ▷ Optimal superposition (c.f. Eq. 3)
if $r < r^*$ **then**
$r^* := r$
$\mathbf{Q}^* := \mathbf{Q}$
$\mathbf{S}^* := \mathbf{S}$
end if
end for
end if
end for
end if
end for
return $\{r^*, \mathbf{Q}^*, \mathbf{S}^*\}$
end procedure

(ACNA) and Neighbour Distance Analysis (NDA), both due to Stukowski [11]. Stukowski uses a Cu-rich 9R precipitate in a BCC-Fe system with 88737 atoms in total as a benchmark to compare a number of structural classification methods. The benchmarked system, which is publicly available [22], is a case where CNA [2], Bond-Angle Analysis [23] and Voronoi-analysis perform badly, and ACNA and NDA perform very well.

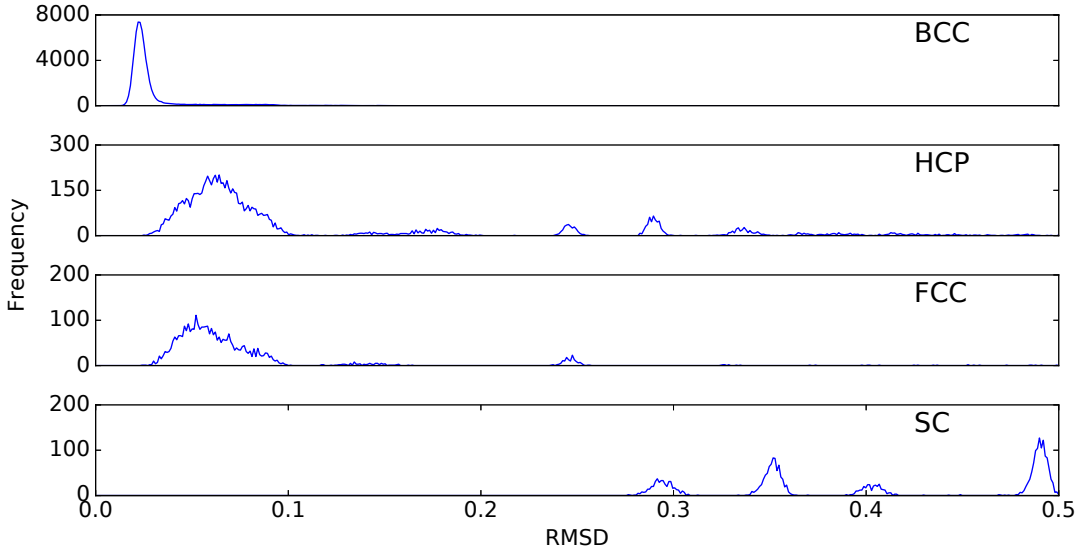

Figure 5. RMSD histograms of BCC, HCP, FCC and SC structures in a Cu-rich 9R precipitate in BCC-Fe system. The BCC atoms have very low RMSD values, which means the local structure is highly ordered. The HCP and FCC structures have both slightly higher RMSD values which are nonetheless crystalline, and very high RMSD spurious identifications. The SC identifications are exclusively spurious and are in fact highly disordered structures.

We have applied our PTM algorithm to the same benchmark system. Figure 5 shows histograms of the RMSD values of the structures identified as BCC, HCP, FCC and SC. We see that a large number of atoms are classified as BCC, FCC and HCP with relatively low RMSD, these are the atoms belonging to locally crystalline areas. In addition, a much lower number of atoms are identified as HCP, FCC or SC but with a much higher RMSD. These are atoms in locally disordered structures, but where the local environment provides a poor match to one of the structural templates. One can choose to eliminate these spurious matches by introducing a cut-off, $\mathrm{RMSD}_{\mathrm{max}}$, for example based on a histogram such as Figure 5.

Table 2 compares the performance of NDA and ACNA with PTM, both with no cut-off ($\mathrm{RMSD}_{\mathrm{max}} = \infty$) and with a sensible cut-off ($\mathrm{RMSD}_{\mathrm{max}} = 0.12$). PTM is capable of indexing even highly distorted structures, which is very useful at high temperatures, but a good cut-off is required to avoid spurious classifications, such as the SC and ICO classifications shown in Table 2. Other than plotting a RMSD-histogram, there is no simple method for selecting $\mathrm{RMSD}_{\mathrm{max}}$, but the cut-off can fortunately be chosen *after* the analysis itself has been performed. The 'correct' value is dependent on the system

Analysis Method	Time (s)	Disordered	SC	FCC	HCP	ICO	BCC
ACNA†	0.66	12789	0	3138	7108	0	65702
NDA†	5.86	13260	0	3453	7825	0	64199
PTM $\mathrm{RMSD}_{\max} = 0.12$	0.82	12900	0	3159	6971	0	65707
PTM $\mathrm{RMSD}_{\max} = \infty$	0.82	7115	2789	3420	8970	8	66435

Table 2. Comparison of of PTM with two other structural analysis methods: Adaptive Common Neighbour Analysis (ACNA) and Neighbour Distance Analysis (NDA). PTM gives similar results to both methods, though with a slightly longer running time than ACNA. The neighbours are ordered by Euclidean distance rather than topologically since the system has been quenched. A good $\mathrm{RMSD}_{\max}$ cut-off is required to avoid false positives. †Both the benchmark data and the implementations of ACNA and NDA are available online [22]. Running times are for analysis only and do not include neighbour-list generation. They were measured on a 2014 MacBook Pro with an Intel Core i7-4770HQ 2.20GHz CPU and 16GB RAM.

being studied, as we will show in section 5.2. For this reason, future publications which use PTM should report the value of $\mathrm{RMSD}_{\max}$ used to ensure reproducibility.

Note that the system here has been quenched, and as a consequence the local structures exhibit little distortion. In this case PTM does not provide any real improvements in structural identification compared to ACNA and NDA. The significant benefits of PTM are to be found in high-temperature systems.

5.2. Copper-Platinum Alloy System at High Temperature

The PTM method is very robust against thermal displacements and strain, and in high-temperature simulations it provides a significant improvement over existing identification methods. We demonstrate this here with a simulated Cu_3Pt alloy sample containing 2.8 million atoms, with periodic boundary conditions. The sample initially contains 30 grains, with randomly selected centres and orientations. The grain volumes have been constructed by computing the Voronoi cells of the grain centres. The atoms within the grains are initially ordered with a perfect $L1_2$ structure. This structure, which is the stable low-temperature phase of this system, can be considered as an FCC structure with three Cu atoms and one Pt atom in the unit cell. It should therefore be identified as FCC by the PTM algorithm.

The sample has first been quenched using the FIRE minimization method [25], and subsequently annealed at 1100K for 820 picoseconds using molecular dynamics with a 5 fs time step. The interatomic potential is the Effective Medium Theory potential [26] and the temperature is controlled with a Berendsen thermostat [27]. After annealing, the sample has been separately heated to 1300K and cooled to 900K, 700K and 500K. Figure 6 compares the performance of ACNA and PTM on the Cu_3Pt systems at 500K and 1100K (both below the melting point temperature) and 1300K (near the melting point temperature). Despite the system at 1300K being in the process of melting from the grain boundaries, PTM correctly identifies the majority of FCC structures in the

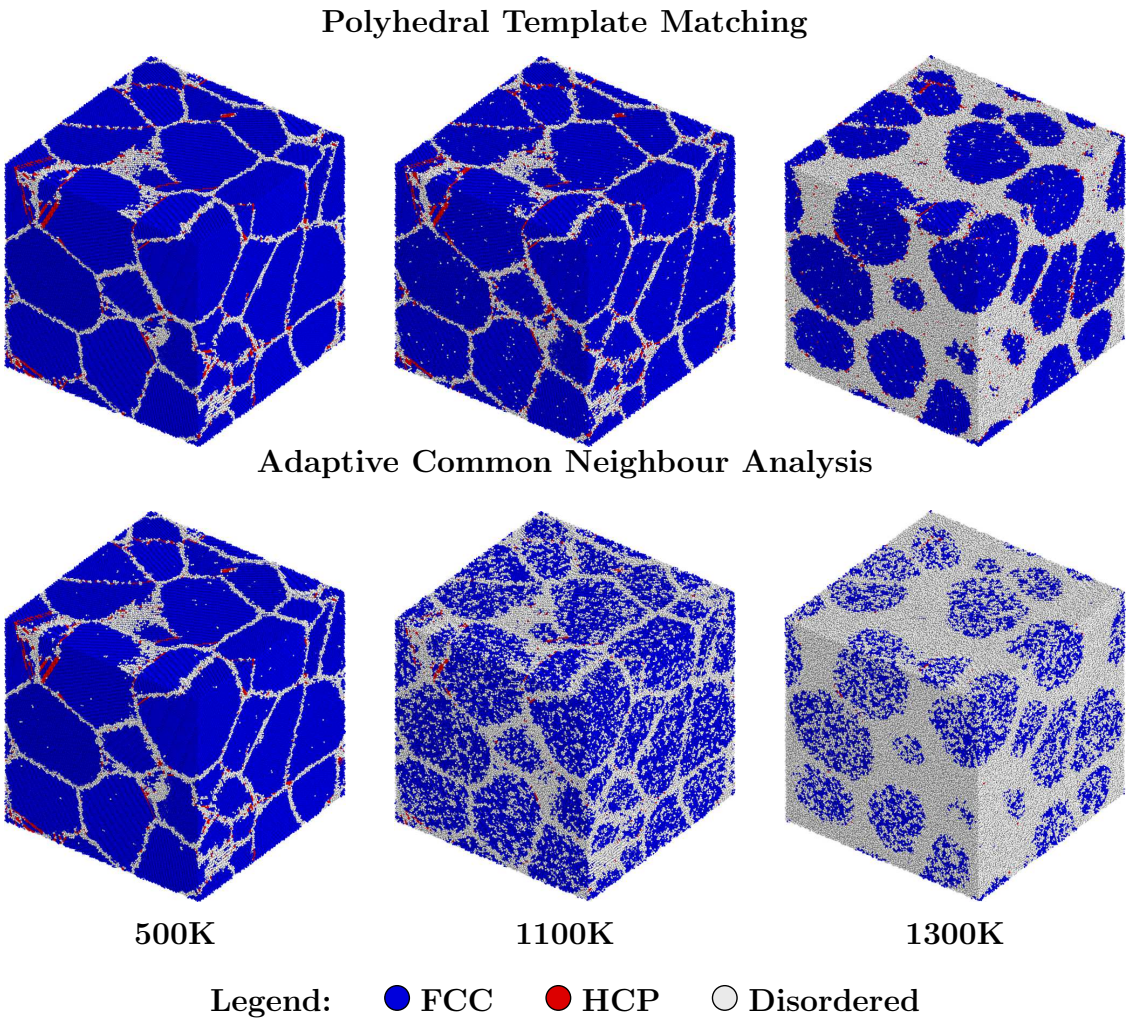

Figure 6. Comparison of PTM with ACNA at three different temperatures on Cu_3Pt systems containing 2.8 million atoms. At 500K PTM offers little improvement over ACNA. With increasing temperature the robustness of PTM is more evident. At 1300K the system is in the process of melting from the grain boundaries. Here, PTM finds twice the number of FCC atoms as ACNA. The RMSD_{max} values used for PTM are 0.11, 0.15 and 0.17 respectively. The choice of RMSD values is motivated by the histograms in Figure 7, but is not critical; using 0.17 for all figures would result in some grain boundary atoms identified as HCP. The SC, ICO and BCC classifications are shown as disordered atoms. The ACNA analyses shown here as well as all renders were performed with OVITO [24].

non-melted volumes. Topological ordering of neighbours improves detection of bulk crystallinity at high temperatures. At 1300K, 9.6% of the atoms identified as FCC with topological ordering were identified as disordered with distance ordering.

Figure 7 shows the effect of temperature on the RMSD distribution of the structures identified as FCC. Higher temperatures lead to more distorted structures. By choosing good values for RMSD_{max}, spurious identifications (in this case SC, ICO and BCC) are significantly reduced. Figure 8 shows the proportion of each structure type identified for the systems shown in Figure 6.

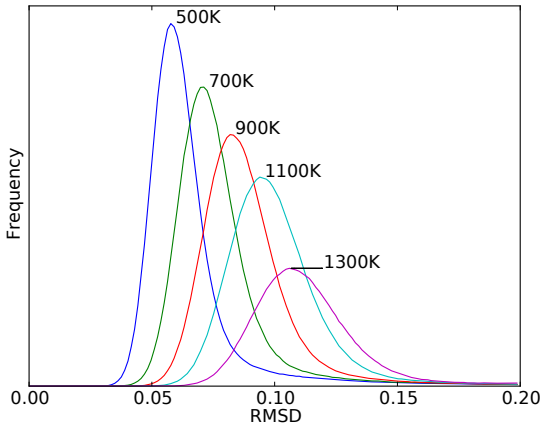

Figure 7. Histograms of RMSD values for the FCC atoms in a polycrystalline Cu_3Pt system at different temperatures. With increasing temperature the thermal displacements result in larger RMSD values.

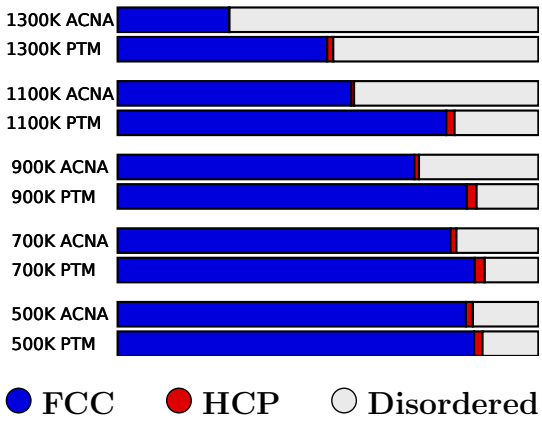

Figure 8. Relative proportions of FCC, HCP and disordered structures found with ACNA and PTM at different temperatures. The SC, ICO and BCC classifications have been counted as disordered.

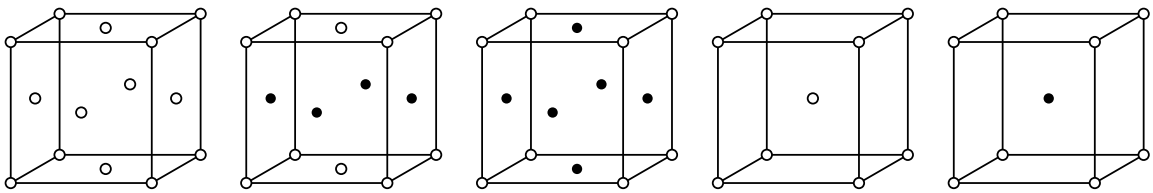

(a) A1 (Cu) **(b)** $L1_0$ (CuAu) **(c)** $L1_2$ (Cu_3Au) **(d)** A2 (W) **(e)** B2 (CsCl)

Figure 9. Lattice structures of the three FCC lattice types (a)-(c), and the two BCC lattice types (d)-(e).

6. Alloy Structures

Since PTM finds the optimal point-to-point correspondences between a set of points and an ideal reference template, we can easily identify alloy structures. A good description of the possible lattice structures of FCC and BCC alloys is given in [28]. Some figures are recreated in Figure 9. Here, we will only consider binary alloys, though multi-element alloys would be a simple extension.

Identification of the A1 and A2 structures is trivial; all nearest neighbours must have the same atom type as the central atom. The B2 structure is equally simple to identify; the types of all the atoms in the first shell of neighbours must be the opposite of the central atom, and all the atoms in the second shell must have the same type as the central atom. For the $L1_0$ and $L1_2$ structures, it is instructive to view the structures from the central atom. This is shown in Figure 10.

Using the point-to-point correspondences, the atom types of a distorted structure

Alloy structure	Atom Type	Neighbours
A1	Cu	3 x 4 coplanar Cu-type (all Cu-type)
A1	Au	3 x 4 coplanar Au-type (all Au-type)
$L1_0$	Cu	2 x 4 coplanar Au-type, 4 coplanar Cu-type
$L1_0$	Au	2 x 4 coplanar Cu-type, 4 coplanar Au-type
$L1_2$	Cu	2 x 4 coplanar Cu-type, 4 coplanar Au-type
$L1_2$	Au	3 x 4 coplanar Cu-type (all Cu-type)

Table 3. Rules for determining FCC binary alloy types using the types of coplanar neighbours.

can be mapped onto the ideal FCC lattice structure. The alloy type can then be determined by examining the types of coplanar neighbours. Table 3 summarizes the rules for determining FCC alloy structures.

To illustrate the method described here, we have determined the alloy structures of the Cu_3Pt system described in section 5.2, both before annealing and after annealing. After the sample grains have been initialized by Voronoi cell construction, with perfect $L1_2$ ordering, the sample has been quenched. The sample has also been quenched, post-annealing, to allow the alloy structures to be compared. The alloy structures of each system are shown in Figure 11. It can be seen that grain boundary migration affects the alloy structure; despite the atoms still having FCC structure after passing through a grain boundary, the $L1_2$ ordering has been lost. Before annealing, less than 0.5% of the FCC atoms were not identified as being in the $L1_2$ structure. After annealing, this percentage had grown to approximately 12%.

7. Lattice Orientations

A consequence of PTM identifying structures by minimizing the RMSD (c.f. Equation (3)) is that the orientation is defined for each atom, at no extra computational cost. Furthermore, due to the optimal point-to-point correspondence being determined, the orientation is robustly determined. Figure 12 shows the lattice orientation of the FCC atoms in the Cu_3Pt system, post-annealing. The non-FCC atoms have not been

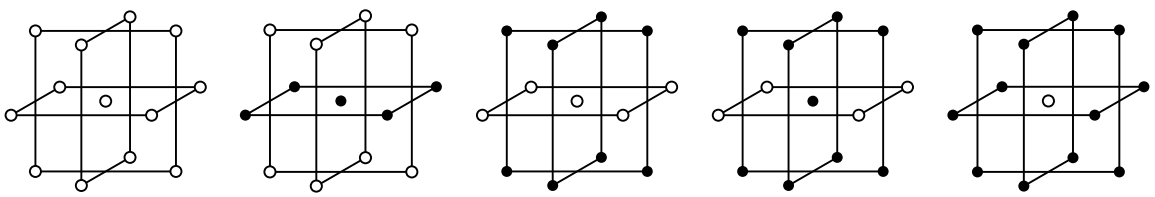

(a) A1 Cu-type **(b)** $L1_0$ Cu-type **(c)** $L1_0$ Au-type **(d)** $L1_2$ Cu-type **(e)** $L1_2$ Au-type

Figure 10. Alloy lattice structures of a central atom and the first shell of nearest neighbours in a FCC lattice, for different central atom types. Different alloy types can be determined by counting the atom types of coplanar neighbours.

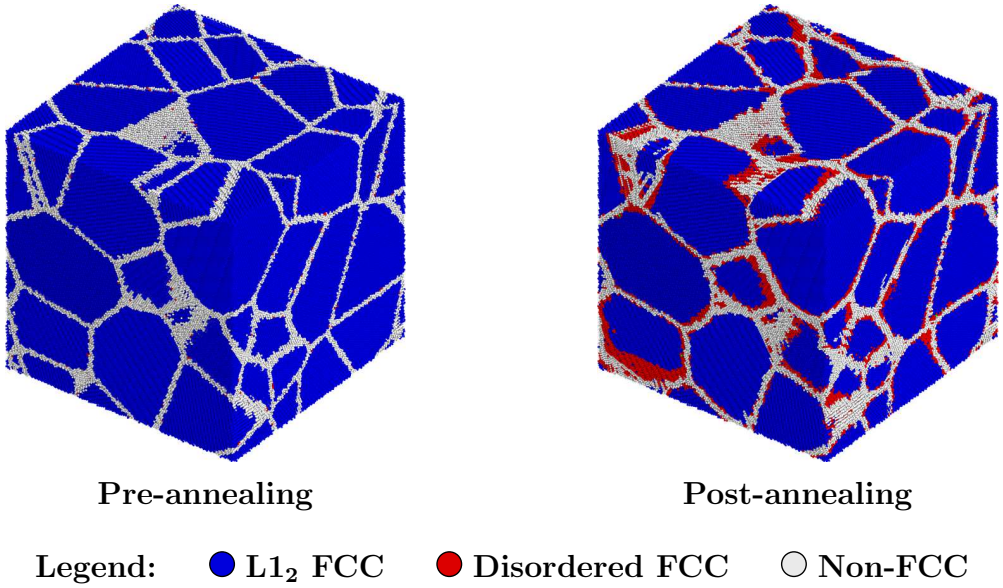

Figure 11. Alloy structures in quenched polycrystalline Cu_3Pt systems. Before annealing (left), the grains have been initialized with perfect $L1_2$ ordering. After annealing at 1100K for 820 picoseconds (right), the atoms which have passed through a grain boundary have disordered FCC alloy structure. The cut-off used is $\mathrm{RMSD}_{\mathrm{max}} = 0.05$.

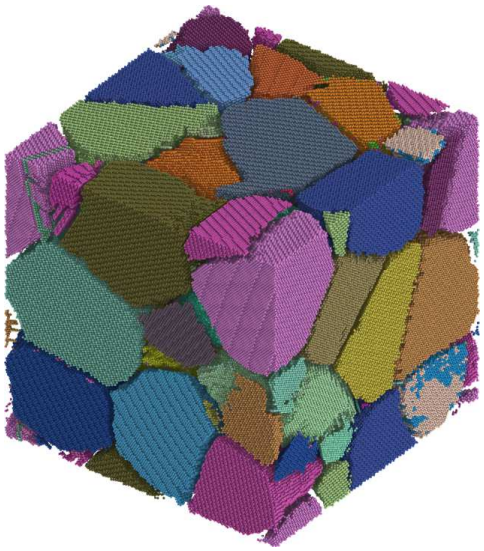

Figure 12. Orientations of FCC atoms in a polycrystalline Cu_3Pt system, with $\mathrm{RMSD}_{\mathrm{max}} = 0.05$.

included in the render. The colours used are obtained by rotation of the orientation of each atom into its fundamental zone and projection into Rodrigues-Frank space. For materials with cubic symmetry, the fundamental zone (in Rodrigues-Frank space) is a truncated octahedron [29], which permits a straightforward mapping into RGB colour space [30, 31]. A drawback of the colour scheme is that grains containing atoms whose fundamental orientations lie close to the faces of the fundamental zone can have very

different colours. This can be seen in the blue/beige grain (lower right) in Figure 12. Nevertheless, the colour scheme conveys the relationship between the grain orientations well.

8. Local Elastic Strain

We have the optimal point-to-point correspondences between the actual atomic positions and the positions of the ideal structure. As such, the local elastic strain is easily obtained, without any reference to an undeformed configuration at the beginning of the simulation. First, we find the deformation gradient using a least-squares fit:

$$r = \min_{\mathbf{A}} ||\mathbf{v}\mathbf{A}^T - \mathbf{w}||_2 \tag{6}$$

where r is the residual term, $\mathbf{v}$ and $\mathbf{w}$ are $3 \times N$ matrices containing the positions of the ideal positions and the optimally permuted actual positions respectively, and the deformation gradient, $\mathbf{A}$, is the affine transformation which minimizes the residual term. The residual term is equivalent to Falk and Langers $D^2_{\min}$ term for identifying local irreversible shear transformations [32]. Prior to fitting the deformation gradient, the ideal and actual positions are translated such that the barycentre of each set lies at the origin, and scaled such that the mean vertex distance is 1, as in Equation (4). Although the deformation gradient obtained is scale-invariant, the scale factor can be used to recover the hydrostatic component. The orientation and elastic strain matrices are obtained via a left-sided polar decomposition of the deformation gradient:

$$\mathbf{P}\mathbf{U} = \mathbf{A} \tag{7}$$

where $\mathbf{U}$ is an orthogonal right-handed matrix (the rotation matrix), and $\mathbf{P}$ is a symmetric matrix (the elastic strain matrix). The choice of a left-sided polar decomposition is arbitrary, but we find the elastic strain in the same frame of reference preferable for comparison of strains across different grains. In the case where $\mathbf{P}$ is not the identity matrix, $\mathbf{U}$ is not the same rotation found by minimizing the RMSD, since the addition of strain means we no longer have a rigid-body transformation. The residual term in Equation (6) could be used to determine the local structure instead of the RMSD, however, the elastic strains in MD simulations are typically less than 5%, and the extra degree of freedom provided by the strain matrix often results in highly-strained spurious structural identifications.

Figure 13 shows the Von-Mises shear strain for the FCC atoms in a $CuPt_3$ system. The Von-Mises shear strain is given by:

$$\varepsilon_{\mathrm{VM}} = \sqrt{\frac{3}{2}\sum_{ij}\mathbf{P}^2_{ij} - \frac{1}{2}\left(\sum_k \mathbf{P}_{kk}\right)^2} \tag{8}$$

where the sums are over the three coordinates. This is the strain equivalent of the Von Mises shear stress, defined in most solid mechanics textbooks [33]. A wedge has been cut out along the length of the system, and a 5% uniaxial strain has been applied transverse

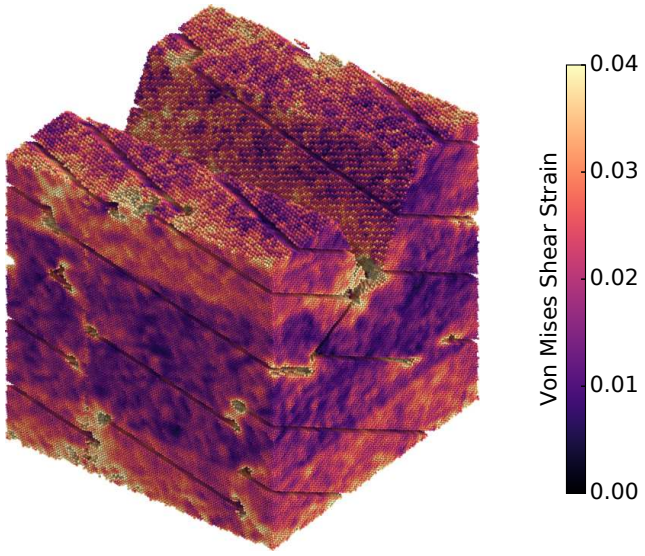

Figure 13. Local Von-Mises shear strain of FCC atoms in a single crystal $CuPt_3$ system, with $RMSD_{max} = 0.02$. Regions of elastic strain occur around the wedge tip, close to surfaces and near dislocations. Strain analysis requires quenching or time-averaging of positions, to avoid strains being dominated by atomic thermal displacements.

to the wedge tip direction. Periodic boundary conditions have been applied along the wedge tip and strain directions. The system was initially a single crystal (to avoid the strain field being dominated by grain boundaries) and has been annealed at 700K for 15000 steps of 5fs to allow stacking faults and dislocations to form. Unfortunately, either quenching or time-averaging of atomic positions is necessary for strain analysis; since the strains are very small, they are dominated by thermal displacements at even moderate temperatures. As such we have quenched the system using the FIRE [25] method. As expected, stress concentrations are seen near dislocation cores and at the tip of the wedge.

9. Implementation

The core components of the PTM algorithm (RMSD optimization, convex hull computation and graph canonization) can be computationally expensive if not implemented carefully. The RMSD optimization is solved using the library provided by D.L. Theobald [15, 34], which is much faster than SVD based calculations. Commonly used convex hull libraries, such as QHull [35] and CGAL [36], implement the QuickHull algorithm (with expected running time $O(n \log n)$) and are designed to achieve good performance on large point sets. On small point sets such as ours theoretical running times are less important to performance than efficient memory allocation, layout and accesses. As such, we have implemented a stack-allocated incremental convex hull algorithm [37] which, despite having a $O(n^2)$ theoretical running time, is 15 times faster than QHull for our point sets. For graph canonization, we have implemented a

single-embedding adaptation of Weinberg's algorithm [19], again with stack allocation. Rather than representing the optimal Weinberg code as a list of edges, we have devised a hash function which uniquely represents all triconnected planar graphs as a 64-bit integer. We have ensured that no hash collisions occur by generating all possible graphs of this type (up to 15 vertices) using the `plantri` program [21]. The deformation gradients in Equation (6) are solved by precomputing a Moore-Penrose pseudoinverse matrix for each structure. This reduces solution of a least-squares problem to matrix multiplication. The polar decomposition of the deformation gradient is computed using a SVD library optimized for 3×3 matrices [38]. For the benchmark system used in section 5.1, the main components account for the following percentages of the total running time: 46% - convex hull construction, 20% - canonical form calculation, 15% - RMSD optimization (including scaling and translation of the points), and 9% - strain calculation. When using Euclidean ordering of neighbouring atoms, the method is capable of indexing over 100,000 atoms per second on a laptop computer using a single core, which is approximately 25% slower than ACNA. Topological neigbour ordering requires the computation of the Voronoi cell of each atom. To do so, we calculate the Delaunay triangulation of a central atom and its 18 nearest neighbours, which allows for up to half the inner shell atoms in FCC and HCP lattices to be wrongly ordered by Euclidean distance. We have implemented the Delaunay triangulation using the parabolic lifting map method of Edelsbrunner and Seidel [39] as it requires fewer intermediate simplices than the Bowyer-Watson [40, 41] algorithm and consequently allows for a stack-allocated implementation. Nonetheless, topological ordering requires significantly more computational effort than Euclidean ordering and increases the running time by a factor of 2. The PTM source code is available online [13].

10. Conclusions

We have presented a computationally efficient yet robust method for identifying the local structure in atomic-scale simulations. The method is based on using the topology of the convex hull formed by the neighbours of an atom to construct a small set of candidate structures, the best of which is chosen based on the root-mean-square deviation of the positions from their expected positions. In case of local crystalline symmetry, the method also identifies the local orientation of the crystalline axes, and optionally the elastic strain tensor. At low to moderate temperatures, the performance of the method is very similar to preexisting methods such as Adaptive Common Neighbour Analysis, but at higher temperatures it is more robust, and can with high reliability identify the structure of a crystalline phase for all temperatures up to the melting point.

Acknowledgments

We would like to thank E. A. Lazar for fruitful discussions of the manuscript. We gratefully acknowledge funding from The Danish National Research Foundation's Center

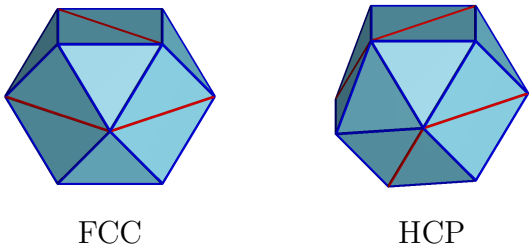

Figure A1. The two different edge types in the convex hull graphs of the FCC and HCP structures; edges which cross a square facet (red) and edges which do not (blue).

for Individual Nanoparticle Functionality (DNRF54).

Appendix A. Symmetrically Inequivalent Triangulations

In section 4 we described how a convex hull can have multiple triangulations of its surface. Here we describe how the symmetrically unique triangulations can be generated, which results in a smaller number of templates to match against a structure.

The SC and ICO structures are simple cases. They both have single convex hull triangulations. Whilst the graph of their convex hulls have 24 and 60 orientation-preserving automorphisms respectively, these automorphisms are symmetrically equivalent. The FCC and HCP cases are slightly more complex. We can generate the unique triangulations by assigning colours to the edges in the convex hull graph. There are two edge lengths in the convex hull; edges which cross a square facet and edges which do not. This is shown in Figure A1. The unique graphs are now the unique edge-colour-preserving and orientation-preserving isomorphisms. The symmetrically unique automorphisms can be determined either by finding the unique RMSD values, or by colouring the facets by type (equilateral triangle or half a square facet) and finding the unique orientation-preserving isomorphisms of the dual graph.

The unique BCC triangulations can be found by colouring the vertices by shell number, and finding the unique vertex-colour-preserving and orientation-preserving isomorphisms. The unique automorphisms can be found either by finding the unique RMSD values, or by colouring the facets according to the number of vertices in each shell and finding the unique orientation-preserving isomorphisms of the dual graph. Whilst we assign colours to edges and vertices during template generation, no colours are assigned during template matching. Colour assignment requires prior knowledge of facet types and, in the BCC case, the neighbour vertex shell numbers. However, the purpose of template matching is to accurately determine the structure. The lack of assumptions about facet types and shell numbers results in robust structural identifications even in highly distorted structures.

References

[1] Kelchner C L, Plimpton S J and Hamilton J C 1998 *Phys. Rev. B* **58** 11085–11088

[2] Honeycutt J D and Andersen H C 1987 *J. Phys. Chem.* **91** 4950–4963
[3] Faken D and Jónsson H 1994 *Comput. Mater. Sci.* **2** 279–286
[4] Schiøtz J and Jacobsen K W 2003 *Science* **301** 1357
[5] Yamakov V, Wolf D, Phillpot S R, Mukherjee A K and Gleiter H 2003 *Philos. Mag. Lett.* **83** 385–393
[6] Brandl C, Derlet P M and Van Swygenhoven H 2011 *Modelling Simul. Mater. Sci. Eng.* **19** 074005
[7] Jónsson H and Andersen H C 1988 *Phys. Rev. Lett.* **60** 2295–2298
[8] Bailey N P, Schiøtz J and Jacobsen K W 2004 *Phys. Rev. B* **69**
[9] Qi Y, Çağin T, Johnson W L and Goddard W A 2001 *J. Chem. Phys.* **115** 385–394
[10] Wells D M, Rossi G, Ferrando R and Palmer R E 2015 *Nanoscale* **7** 6498–6503
[11] Stukowski A 2012 *Modelling Simul. Mater. Sci. Eng.* **20** 045021
[12] Lazar E A, Han J and Srolovitz D J 2015 *Proc. Natl. Acad. Sci. USA* **112** E5769–E5776
[13] Larsen P M, Schmidt S and Schiøtz J 2016 Polyhedral Template Matching - Source Code Repository `http://github.com/pmla/polyhedral-template-matching`
[14] Horn B K P 1987 *J. Opt. Soc. Am. A: Opt. Image Sci. Vis.* **4** 629–642
[15] Theobald D L 2005 *Acta Crystallogr. A* **61** 478–480
[16] Steinitz E 1916 *Polyeder und Raumeinteilungen* (Teubner)
[17] Whitney H 1933 *Am. J. Math.* **55** 231
[18] Tutte W T 1963 *Proc. London Math. Soc.* **s3-13** 743–767
[19] Weinberg L 1966 *IEEE Trans. Circuit Theory* **13** 142–148
[20] Kukluk J P, Holder L B and Cook D J 2004 *J. Graph Algorithms Appl.* **8** 313–356
[21] Brinkmann G and McKay B D 2007 *MATCH Commun. Math. Comput. Chem.* **58** 323–357
[22] Stukowski A 2012 Atomic structure analysis algorithms `http://asa.ovito.org`
[23] Ackland G J and Jones A P 2006 *Phys. Rev. B* **73** 054104
[24] Stukowski A 2010 *Modelling Simul. Mater. Sci. Eng.* **18** 015012
[25] Bitzek E, Koskinen P, Gähler F, Moseler M and Gumbsch P 2006 *Phys. Rev. Lett.* **97** 170201
[26] Jacobsen K W, Stoltze P and Nørskov J K 1996 *Surf. Sci.* **366** 394–402
[27] Berendsen H J C, Postma J P M, van Gunsteren W F, DiNola A and Haak J R 1984 *J. Chem. Phys.* **81** 3684–3690
[28] Laughlin D E 1988 Long-range order in alloys *Encyclopedia of Materials Science and Engineering, Supp. Vol. 1.* (Pergamon Press) pp 263–268
[29] He Y and Jonas J J 2008 Rodrigues-Frank Spaces for Misorientations and Orientation Relationships between Crystals of Any Two Crystallographic Point Groups *Applications of Texture Analysis Ceramic Transactions* (Hoboken, NJ, USA: John Wiley & Sons, Inc.) pp 269–280
[30] Albou A, Driver J H and Maurice C 2010 *Acta Mater.* **58** 3022–3034
[31] Van Boxel S, Seefeldt M, Verlinden B and van Houtte P 2005 *J. Microsc.* **218** 104–114
[32] Falk M L and Langer J S 1998 *Phys. Rev. E* **57** 7192–7205
[33] Budynas R G and Nisbett J K 2011 *Shigley's Mechanical Engineering Design* McGraw-Hill series in mechanical engineering (McGraw-Hill)
[34] Liu P, Agrafiotis D K and Theobald D L 2010 *J. Comput. Chem.* 1561–1563
[35] Barber C B, Dobkin D P and Huhdanpaa H 1996 *ACM Trans. Math. Software* **22** 469–483
[36] Hert S and Schirra S 2013 3D Convex Hulls *CGAL User and Reference Manual* (CGAL Editorial Board) URL `http://doc.cgal.org/4.2/CGAL.CGAL/html/packages.html#PkgConvexHull3Summary`
[37] O'Rourke J 1998 *Computational geometry in C* (Cambridge University Press)
[38] McAdams A, Selle A, Tamstorf R, Teran J and Sifakis E 2011 Computing the singular value decomposition of 3×3 matrices with minimal branching and elementary floating point operations Tech. Rep. 1690 University of Wisconsin-Madison
[39] Edelsbrunner H and Seidel R 1986 *Discrete Comput. Geom.* **1** 25–44
[40] Bowyer A 1981 *Comput. J.* **24** 162–166
[41] Watson D F 1981 *Comput. J.* **24** 167–172